

TASEP IN ANY WEYL GROUP

ERIK AAS

Abstract. We investigate a Markov chain defined by Thomas Lam [6], which generalizes the multi-type TASEP on a ring to any Weyl group. For groups of type C we define an analogue of the multiline queues of Ferrari and Martin (which compute the stationary distribution for the classical TASEP). While our construction does not suffice for finding the stationary distribution, the construction gives the stationary distribution of a certain projection of Lam's chain. Also, our approach is incremental, in the sense that the construction appears to fit into a pattern of 'conjugation matrices', which remains to be fully worked out. Finally, we prove a theorem for the classical TASEP which fits into the picture of viewing TASEP in a permutation-free way.

1. INTRODUCTION

Let W be any finite Weyl group. In [6], Lam defined an interesting Markov chain Θ'_W on W whose stationary distribution is connected to the study of reduced expressions in the corresponding affine group. It turns out [1] that for W of type A, Θ'_W is the well-known multi-type TASEP on a ring, studied for some time by probabilists. In particular its stationary distribution has a very elegant description in terms of the *multi-line queues* of Ferrari and Martin [4]. One could hope for a similar description of the stationary distribution of Θ'_W for general W . This appears to be difficult. Lam also defined a weighted variant Θ_W whose stationary distribution seems more amenable to analysis (though its connection to reduced words is unclear).

- For each W , we identify certain projections of Θ_W , which generalize the familiar operation of merging different particle classes in the TASEP. We state this in its most general form in Section 5.
- For W of type C, we compute the stationary distribution of one of these projections. We also show how to 'invert' some of the projections. That is, we give a way to simulate one projection Θ_1 in terms of a further projection Θ_2 of it, together with independent randomness. One of these inversions can be seen as an analogue of the multi-line queues of Ferrari and Martin (see Theorem 3.1).
- We give an equivalent description of the classical (type A) TASEP which seemingly can be formulated purely in terms of geometrical notions (i.e. independently of the permutation representation of the Weyl group of type A). Though it has some interest in itself, we do not manage to generalize it to other Weyl groups.
- In the final section we make some concluding remarks and pose several questions.




FIGURE 1. A state.

The paper is structured as follows. In Section 2, we describe Lam's chain in the special case we are most interested in, the type C TASEP, and explain the notion of projections of Markov chains in this context. In Sections 3 and 4 we prove two theorems partially describing the stationary distribution of Θ_W for W of type C. In Section 5 we generalize some of the results in the preceding sections to arbitrary Weyl groups. In Section 6 we prove a theorem on parallel update rules for the classical TASEP on a ring, which coincides with Θ_W for W of type A. Finally in Section 7 we pose some questions which we have no satisfactory answer to.

2. TYPE C

Fix $n \geq 2$. Throughout this section we fix W to be the Weyl group of type C and rank n . Rather than work through Lam's definition of $\Theta = \Theta_W$ (given in Section 5), we give another definition which is easily checked to be equivalent (using the well-known permutation presentation of W , see [3]).

A *state* in Θ is an assignment of labeled circles to the sites of a cycle of length $2n$ drawn as in Figure 1 (in the case $n = 10$).

We refer to the circles in the diagram as *particles*, and the numbers in them as their corresponding *classes*. No two particles are allowed to occupy the same column, though columns are allowed to be empty (such as column number 8 from the left in the example).

There is an exponential bell with rate 1 at each site. When a bell is activated, the particle at that position (if there is one) either (i) jumps to the next site s counter-clockwise if that position is empty (ii) trades place with the particle at s if that particle has higher class (iii) does nothing if the particle at s has lower class. Moreover, the bell in the upper line in column i and the bell in the lower line in column $i + 1$ trigger each other, for each $1 \leq i < n$. This ensures that no two particles occupy the same column at any time. We may thus think of the pair of bells as a single bell with rate 2. There are then n bells with rate 2 in the middle and two bells with rate 1 in the leftmost and rightmost column. We denote the action of a bell at site i in the lower line and site $i + 1$ in the upper line by σ_i , for $1 \leq i < n$. The bell at the site furthest down to the right similarly defines σ_n and the bell furthest up to the left defines σ_0 .

This defines our Markov chain Θ . Clearly the number m_i of particles of each class i is conserved by the dynamics. We refer to the vector $\mathbf{m} = (m_1, m_2, \dots)$ as the *type* of the state.

When there is one particle of each class, and each column is occupied by one particle, states are in one-to-one correspondence with signed permutations w of $[n]$ – there is a particle of class i in the upper row in column j if $w(j) = +i$ and in the lower row if $w(j) = -i$. In the general case, states correspond to left cosets of the group of signed permutations – we will explain this in more detail later.

For $J \subseteq [n]$ we define a type \mathbf{m}_J as follows. Start with a state of type $(1, 1, \dots, 1)$. If $j \in J - \{n\}$, identify the particle classes j and $j+1$. If $n \in J$, remove all particles of class n . Now renumber the remaining particle classes with integers $1, 2, \dots$. The resulting state has a type which we define to be \mathbf{m}_J . For example, the states in Figures 1 and 2 have type $\mathbf{m}_{\{1,2,4,6,10\}}$ and $\mathbf{m}_{\{1,2,4,6\}}$ respectively. Denote the set of all states of type \mathbf{m}_J by Ω_J .

Clearly the \mathbf{m}_J enumerate all the interesting variations of types of states. Using the identifications above we similarly obtain a projection map φ_i from words of type \mathbf{m}_J to words of type $\mathbf{m}_{J \cup \{i\}}$ for each J such that $i \notin J$. We denote the restriction of Θ to states of type \mathbf{m}_J by Θ_J . It is easy to see that Θ_J is aperiodic and irreducible. Thus it has a unique stationary distribution π_J .

Suppose (i, J, J') satisfies $i \notin J \subseteq [n]$, $J' = J \cup \{i\}$. We should think of (i, J, J') as a covering relation in the boolean poset of all subsets of $[n]$, and call such a triple a *link*. Of course any two of i, J and J' determine the third, but the terminology of links turns out to be convenient.

Proposition 2.1. *Suppose (X_1, X_2, \dots) is a random walk in Θ_J , and (i, J, J') is a link. Then $(\varphi_i(X_1), \varphi_i(X_2), \dots)$ is a random walk in $\Theta_{J'}$.*

Definition 2.2. *For a Markov chain Θ with state space Ω , we define its $\Omega \times \Omega$ transition matrix M_Θ by letting*

$$M_\Theta(u, v) = \sum_{\gamma: u \rightarrow v} \text{rate}(\gamma),$$

the sum of the rates of all transitions γ from state u to state v .

Write M_J for M_{Θ_J} . Thus the stationary distribution π_J is an eigenvector of M_J , with eigenvalue $2n$.

Suppose (i, J, J') is a link, and define a $\Omega_{J'} \times \Omega_J$ matrix $D = D_{i,J}$ by letting $D(v, u)$ be 1 if $v = \varphi_i(u)$, and 0 otherwise. Proposition 2.1 can be strengthened as follows.

Proposition 2.3. *Suppose (i, J, J') is a link. Then*

$$D_{i,J} M_J = M_{J'} D_{i,J}.$$

Proof. Choose $(v, u) \in \Omega_{J'} \times \Omega_J$. The (v, u) entry of the left hand side counts the number of $j \in [0, n]$ (where 0 and n are counted with weight 1 and the others with weight 2) such that $v = \varphi_i(\sigma_j u)$. Similarly, the (v, u) entry of the right hand side counts the number of j' such that $v = \sigma_{j'}(\varphi_i(u))$. These two counts can be matched to each other simply by taking $j = j'$. \square

Proposition 2.3 is stronger than Proposition 2.1 in the sense that for *any* eigenvector v of M_J with eigenvalue λ , we have $\lambda D_{i,J} v = D_{i,J} M_J v = M_{J'}(D_{i,J} v)$, so that $D_{i,J} v$ is an eigenvector of $M_{J'}$ if it is non-zero. Since π_J spans a one-dimensional eigenspace of M_J for each J , $D_{i,J} \pi_J$ must be a scalar multiple of $\pi_{J'}$. Though we are only interested in the stationary distribution of the Θ_J , considering *all* the eigenvectors will help us to compute the stationary distribution π_J .¹

Of course, Proposition 2.3 appears to have no practical use – the chain Θ_J is more complicated than the chain $\Theta_{J'}$. Therefore, the following fact from linear algebra is quite a revelation.

¹In [2], the authors investigate the eigenvalues and eigenvectors for the TASEP on a ring, ie. our chain when W is of type A.




FIGURE 2. The result of applying $\tau_{(+,-,-,+,+,-,-,-,-)}$ to the state in Figure 1.

Proposition 2.4. *Suppose A, B are matrices. If there is a matrix D of full rank such that*

$$DA = BD,$$

then there is a matrix U of full rank such that

$$AU = UB.$$

So the mere existence of the *projection matrix* $D_{i,J}$ implies the existence of some *conjugation matrix* $U_{i,J}$ such that

$$M_J U_{i,J} = U_{i,J} M_{J'},$$

which would allow us to compute π_J from $\pi_{J'}$ by $\pi_J = U_{i,J} \pi_{J'}$! Of course there is no guarantee that there will be a "simple" matrix $U_{i,J}$ (eg. with small positive integer entries). However, in the coming sections we will identify such $U_{i,J}$ for some links (i, J, J') . It would be very interesting if these $U_{i,J}$ could be defined at the generality of Proposition 2.4 – i.e. as a function of $(M_J, M_{J'}, D_{i,J})$. The two cases we consider will correspond to adding/removing the particles of largest class, and the case of only two classes of particles.

3. ADDING PARTICLES OF LARGEST CLASS

In this section we construct a conjugation matrix $U = U_{i,J}$ for links (i, J, J') with $i = n$ and $J \subseteq [n-1]$ arbitrary. To describe U , take a state u of type $\mathbf{m}_{J'}$, and let $\vartheta \in \{+, -\}^n$. We will produce a new state $v = \tau_{\vartheta} u$, of type \mathbf{m}_J . Column j in v will be occupied in the upper line if $\vartheta_j = +$ and in the lower line otherwise (in particular, there will be no empty columns in v). In the start we refer to these sites as Not Yet Occupied. Go through the particles in u in any order such that particles with small class come before particles with larger class. The order of particles of same class does not matter. Now, when considering a particle p at a site s in u , find the first NYO site in v , going *clockwise* from s . Put p at this site. Now s has been occupied. When all particles in u have been processed, fill the remaining NYO sites in v with particles of a new largest class.

In Figure 2, we have applied τ_{ϑ} , where $\vartheta = (+, -, -, +, +, -, -, -, -, -)$, to the state in Figure 1. We define U by letting its (u, v) entry be the number of $\vartheta \in \{+, -\}^n$ such that $v = \tau_{\vartheta} u$.

Theorem 3.1. *In the notation above, we have*

$$U M_{J'} = M_J U.$$

Proof. Fix $u \in \Omega_{J'}$. We need to show that the number (counted with the same weights as before) of (j, ϑ) such that $v = \sigma_j \tau_{\vartheta} u$ equals the number of (j', ϑ') such that $v = \tau_{\vartheta'} \sigma_{j'} u$, for each $v \in \Omega_J$.




FIGURE 3. In the top left corner we have drawn a state u , and indicated the actions of $\sigma_{j'}$ and τ_{ϑ} . Below and to the right are $\sigma_{j'}u$ and $\tau_{\vartheta}u$ with $\tau_{\vartheta'}$ and σ_j indicated respectively. The claim is that the state in the bottom right is both $\sigma_j\tau_{\vartheta}u$ and $\tau_{\vartheta'}\sigma_{j'}u$. We have only indicated one possible result $\tau_{\vartheta}u$; it could be that the position labeled 3 is replaced by a particle with smaller class. However, since there are no particles of class < 1 we are certain that the position labeled 1 is labeled so in any circumstance. It is easy to argue for each particle that it is going to end up in the same place in both $\tau_{\vartheta'}\sigma_{j'}u$ and $\sigma_j\tau_{\vartheta}u$.

Some pairs (j, ϑ) have a natural corresponding pair (j', ϑ') with $j' = j$, with same weight and such that $\tau_{\vartheta}\sigma_j u = \sigma_{j'}\tau_{\vartheta'} u$, as follows.

- if $0 < j < n$ and $\vartheta_j = \vartheta_{j+1}$, let $\vartheta' = \vartheta$.
- if $0 < j < n$ and $(\vartheta_j, \vartheta_{j+1}) = (-, +)$, let $\vartheta' = (\vartheta_1, \dots, \vartheta_{j-1}, \vartheta_{j+1}, \vartheta_j, \vartheta_{j+2}, \dots, \vartheta_n)$.
- if $j = 0$ and $\vartheta_1 = -$, let $\vartheta' = (+, \vartheta_2, \dots, \vartheta_n)$.
- if $j = n$ and $\vartheta_n = +$, let $\vartheta' = (\vartheta_1, \dots, \vartheta_{n-1}, -)$.

We illustrate the first of these in Figure 3.

Thus we need to show that the set S of remaining pairs (j, ϑ) , ie. those satisfying

- $0 < j < n$, $(\vartheta_j, \vartheta_{j+1}) = (-, +)$,
- $j = 0$, $\vartheta_1 = +$, or
- $j = n$, $\vartheta_n = -$,

have the same effect as the set S' of remaining pairs (j', ϑ') , ie. those satisfying

- $0 < j < n$, $(\vartheta_j, \vartheta_{j+1}) = (+, -)$,
- $j = 0$, $\vartheta'_1 = -$, or
- $j = n$, $\vartheta'_n = +$.

It is easy to see that for all $(j, \vartheta) \in S$, the state $\tau_{\vartheta}\sigma_j u$ is the same, and equal to $\sigma_{j'}\tau_{\vartheta'} u$ for all $(j', \vartheta') \in S'$. Thus it suffices to show that their (weighted) counts are the same! This is easily done by considering separately the cases $(\vartheta_1, \vartheta_n) = (+, +), (+, -), (-, +), (-, -)$ and similarly for $(\vartheta'_1, \vartheta'_n)$.

□

Thus, if u is distributed according to $\pi_{J'}$ and $\vartheta \in \{+, -\}^n$ is chosen independently and uniformly at random, then $\tau_{\vartheta}u$ will be distributed according to π_J .

Corollary 3.2. *The Markov chain on Ω_J obtained by at each time step applying a random τ_{ϑ} has the same stationary distribution as Θ_J .*

Proof. In the proof of Theorem 3.1, we never used the fact that there was some empty column in the state u . Thus the same proof shows that $M_J U = U M_J$ where $U(u, v) = 1$ if $v = \tau_{\vartheta}u$ for some ϑ and 0 otherwise, for $(u, v) \in \Omega_J \times \Omega_J$. Thus U maps π_J onto a constant multiple of itself. □

4. TWO PARTICLE CLASSES

In this section we construct $U = U_{t,J}$ for the case $J = [n] - \{t\}$, $t \neq n$ (note that the case $t = n$ is trivial). In this case U will only have one column, so this is equivalent to describing the stationary state of Θ_J . We will not phrase the result in terms of U .

The chain Θ_J has strong similarities with the 3-type TASEP on a ring, and with the chain studied in [5]. Indeed we will show that the stationary distribution satisfies similar recursion relations to these two chains (only the initial data will be different), and the proof is very similar to that of [5].

When $J = [n] - \{t\}$, we are considering states with t particles of class 1. It will be more convenient to write states as words, as follows. A state u is described by a word $w_1, \dots, w_n \in \{-1, 0, 1\}^n$, where $w_i = 1$ if there is a particle (with the only class 1) in the upper line in column i , $w_i = -1$ if there is one in the lower line, and $w_i = 0$ if both lines are empty in this column. We sometimes write $-x = \bar{x}$ for $x = 1, 0, \bar{1}$. It would be more natural to write 1, 2, 3 for 1, 0, $\bar{1}$, though we will not do this in order to be consistent with previous sections.

In this notation, Θ_J becomes a chain on words, where any subword 01, $\bar{1}1$, $\bar{1}0$ turns into 10, $1\bar{1}$, $0\bar{1}$ respectively at rate 2, a $\bar{1}$ on the right end turns into 1 at rate 1, and a 1 at the left end turns into a $\bar{1}$ at rate 1.

Our analysis becomes more transparent if we temporarily generalize Θ_J so that a subword 01 turns into 10 at rate a , $\bar{1}1$ into $1\bar{1}$ at rate b , $\bar{1}0$ into $0\bar{1}$ at rate c , $\bar{1}$ on the right end into 1 at rate d , and 1 at the left end into $\bar{1}$ at rate e , for indeterminates a, b, c, d, e .

Furthermore, we will not consider the length n of the words to be fixed anymore; we want to describe the value of stationary distribution at *any* word in $\{-1, 0, 1\}$ of length at least 2.

Thus our case of interest is (a, b, c, d, e) proportional to $(2, 2, 2, 1, 1)$. To get conventions right, the most convenient choice will be $(a, b, c, d, e) = (1, 1, 1, \frac{1}{2}, \frac{1}{2})$.

Definition 4.1. *For each word u (of any length ≥ 2) we define a Laurent polynomial $[u]$ in a, b, c, d, e as follows. Let v, w be any words (possibly empty). Then*

$$(1) \quad [v01w] = [v0w]/a,$$

$$(2) \quad [v\bar{1}1w] = ([v\bar{1}w] + [v1w])/b,$$

$$(3) \quad [v\bar{1}0w] = [v0w]/c,$$

$$(4) \quad [v\bar{1}] = [v]/d,$$

$$(5) \quad [1v] = [v]/e.$$

Moreover, if u consists of 0s only, then $[u] = 1$.

To prove that this defines $[u]$ in a unique way, we need to show that when expanding according to the recursions above we always arrive at the same result $t \cdot [0^s]$, where t is some Laurent polynomial and s is the number of 0's in u (which is clearly conserved by the recursions - in fact we could have made any set of choices for $[], [0], [00], [000], \dots$ rather than setting them all equal to 1). This is easy by induction, using the following induction hypothesis:

For words of length $\leq n$, taking any definition of $[u]$ for each u , all equations between brackets of words of length $\leq n$ are satisfied.

Here is an interesting consequence of the definition.

Proposition 4.2. *If $u = v_1 0 v_2 0 \dots 0 v_r$ where v_1, \dots, v_r are any words in $\{-1, 0, 1\}$, then $[u] = [v_1 0][0 v_2 0] \dots [0 v_r]$.*

Theorem 4.3. *For any n and t , the stationary distribution of Θ_J , where $J = [n] - \{t\}$, evaluated at the state u is proportional to $[u]$.*

Proof. It suffices to prove that the numbers $[u]$ satisfy the equilibrium equation, which reads

$$(6) \quad 0 = T_0 + \sum_{0 < i < n} T_i + T_n,$$

where

$$T_0 = \text{wt}(u_1 \rightarrow \bar{u}_1)[u_1 \dots u_n] - \text{wt}(\bar{u}_1 \rightarrow u_1)[\bar{u}_1 u_2 \dots u_n],$$

$$T_i = \text{wt}(u_i u_{i+1} \rightarrow u_{i+1} u_i)[u_1 \dots u_n] - \text{wt}(u_{i+1} u_i \rightarrow u_i u_{i+1})[u_1 \dots u_{i-1} u_{i+1} u_i u_{i+2} \dots u_n],$$

for $0 < i < n$ and

$$T_n = \text{wt}(u_n \rightarrow \bar{u}_n)[u_1 \dots u_n] - \text{wt}(\bar{u}_n \rightarrow u_n)[u_1 \dots u_{n-1} \bar{u}_n].$$

Here $\text{wt}(01 \rightarrow 10) = a$, $\text{wt}(1\bar{1} \rightarrow \bar{1}1) = 0$, $\text{wt}(1 \rightarrow \bar{1}) = e$ at the left end, et.c. A case by case analysis (using 1-5) shows that

$$T_0 = a_{u_1}[u_2 \dots u_n],$$

$$T_i = a_{u_i}[u_1 \dots u_{i-1} u_{i+1} \dots u_n] - a_{u_{i+1}}[u_1 \dots u_i u_{i+2} \dots u_n]$$

for $0 < i < n$ and

$$T_n = -a_{u_n}[u_1 \dots u_{n-1}]$$

where $(a_1, a_0, a_{\bar{1}}) = (1, 0, -1)$.

This turns the right hand side of (6) into a telescoping sum with value 0 (each term $a_{u_i}[\dots \hat{u}_i \dots]$ occurs twice - once with each sign), so the equilibrium equation is satisfied. □

For the case we are interested in, $(a, b, c, d, e) = (1, 1, 1, \frac{1}{2}, \frac{1}{2})$, we obtain

Corollary 4.4. *For each n, t , choose $\alpha_{n,t} > 0$ such that the minimum of the numbers $n_u = \alpha_{n,t} \pi_J(u)$ for $u \in \Omega_{[n]-t}$ is 1. Then*

- Each n_u is a positive integer.
- $n_u = 1$ if and only if $u = \bar{1}^i 0^j 1^k$ for some i, j, k (satisfying $i + k = t$, $i + j + k = n$).

- $n_u \leq 2^t$ for all u , with equality if and only if $u = 1^i 0^j \bar{1}^k$ for some i, j, k .
- For any words u, v, w in $\{-1, 0, 1\}$, $n_{u0v0w} = n_{u0} \cdot n_{0v0} \cdot n_{0w}$.

5. THE GENERAL CASE

In this section, we explain how much of the analysis carries over to general Weyl groups. The short answer is that the general setup can be formulated for a general Weyl group, but the lack of a useful permutation representation makes it hard to describe the conjugation matrices in an easy way.

Fix a root system Φ of rank n in \mathbb{R}^N (for some N) with a simple system $\alpha_1, \dots, \alpha_n$. Let α_0 be the highest root with respect to this choice of simple system, and write $\alpha_0 \alpha_0 = \sum_{i=1}^n a_i \alpha_i$ where the a_i are nonnegative integers with $a_0 = 1$. For $i \in [0, n]$, we let t_i be the reflection in the hyperplane orthogonal to α_i . Then t_1, \dots, t_n generate the Weyl group of Φ . We denote its length function by $\ell(\cdot)$. For an element $w \in W$ and $i \neq 0$, we define $\sigma_i(w)$ to be equal to w if $\ell(wt_i) > \ell(w)$ and equal to wt_i otherwise. For $i = 0$, we let $\sigma_i(w) = w$ if $\ell(wt_0) < \ell(w)$ and wt_i otherwise.

The following is Lam's original definition of Θ .

Definition 5.1. *The Markov chain Θ has W as state space, and the outgoing transition from any state $w \in W$ are all $w \rightarrow \sigma_i(w)$, $i \in [0, n]$. Transitions corresponding to σ_i have rate a_i .*

For type A, we have $(a_0, \dots, a_n) = (1, \dots, 1)$ and the chain Θ' in Section 1 coincides with Θ in this case. For type C we have $(a_0, a_1, \dots, a_{n-1}, a_n) = (1, 2, \dots, 2, 1)$. The chain for type C is equivalent to the chain defined in Section 1.

It appears that the above choice of rates (a_0, \dots, a_n) is essentially the only one that gives a nice stationary distribution for Θ (eg. such that the probability of any state divided by the probability of a least likely state is an integer, and these integers are not 'too large'). I have no good explanation for this experimental fact; however, see Remark 5 in [6].

We now show how to generalize Propositions 2.1 and 2.3 to the general setting. For $J \subseteq [n]$, let W_J be the subgroup of W generated by $\{t_j : j \in J\}$. Note that we do not allow $0 \in J$.

Proposition 5.2. *Suppose $w, w' \in W$ satisfy $W_J w = W_J w'$. Then, for any $i \in [0, n]$, we have $W_J(\sigma_i(w)) = W_J(\sigma_i(w'))$.*

Proof. This is clear if $\ell(wt_i) > \ell(w)$ and $\ell(w't_i) > \ell(w')$, or if $\ell(wt_i) < \ell(w)$ and $\ell(w't_i) < \ell(w')$.

We can thus focus on the case when $\ell(wt_i) > \ell(w)$ and $\ell(w't_i) < \ell(w')$, the fourth case then follows by symmetry in w, w' . It suffices to prove that $wt_i w^{-1} \in W_J$ since this will prove that $W_J \sigma_i w' = W_J w' t_i = W_J w t_i = W_J w = W_J \sigma_i w$ if $i \neq 0$ and $W_J \sigma_i w' = W_J w' = W_J w = W_J w t_i = W_J \sigma_i w$ if $i = 0$. Since each W_J has a minimal element (Corollary 2.4.5 in [3]), we can reduce to the case when $w < w'$ in Bruhat order (the case $w' < w$ is similar).

Then (see the mirrored version of Proposition 2.4.4 in [3]), w has a reduced expression

$$s_1 \dots s_r,$$

and w' has a reduced expression

$$s'_1 \dots s'_l$$

where $s'_i \in W_J$.

The conditions imply (Corollary 1.4.4 in [3]) that $t_i = s_r \dots s_1 s'_l \dots s'_k \dots s'_l s_1 \dots s_r$ for some k and thus $ut_i u^{-1} = s'_l \dots s'_k \dots s'_l \in W_J$. \square

Thus all the subgroup inclusion $W_{J'} \leq W_J$ (with $J' \subseteq J$) induce projections $D_{J,J'} : \Theta_J \rightarrow \Theta_{J'}$ in the same way as in Section 2. Explicitly, let $D_{i,J}(W_J u, W_{J'} v) = 1$ if $W_{J'} u = W_{J'} v$ and 0 otherwise. Then $D_{i,J} M_J = M_{J'} D_{i,J}$. We should thus be able to expect there to be corresponding conjugation matrices $U_{J,J'}$. We have not succeeded in finding a general description of these (indeed, not even for groups of type C), though there appears to be 'surprisingly simple' such U in several cases.

6. THE k -TASEP

In this section, we prove a theorem about the TASEP on a ring (i.e. the type A case of the chains considered above), which can be formulated independently of the permutation representation. However the most obvious generalizations of the theorem do not appear to hold for a general Weyl group.

We now define the classical multi-type TASEP on a ring. We will include indeterminates x_1, x_2, \dots in the definition as in [7]. This is more general than the chain Θ_W considered earlier, but reduces to it on letting $x_1 = x_2 = \dots = 1$.

The state space is the set of words of length n in the alphabet $\{1, 2, \dots\}$. We consider the words to be cyclic, so that the letter to the left of w_1 is $w_0 = w_n$ (all indices are taken modulo n). For $i \in [n]$ we define $\sigma_i(w)$ as the result of sorting the two letters w_i and w_{i-1} . Thus if the letters already satisfy $w_{i-1} \leq w_i$, nothing happens, and otherwise they swap positions.

The outgoing transitions from a general state (word) u are all $u \rightarrow \sigma_i u$, $i \in [n]$, where transitions corresponding to σ_i have rate x_{u_i} (note that this definition is *not* symmetric in (u_{i-1}, u_i)).

For a proper subset $S \subsetneq [n]$ we would like to define σ_S as the composition of all σ_j , $j \in S$. To do this we need to specify in which order non-commuting pairs of σ_j 's should be taken. Note that σ_j and σ_k commute whenever $|k - j| > 1$ modulo n . We use the convention that σ_{j-1} is taken before σ_j . Thus, for example, if $n = 7$, $\sigma_{\{1,2,4,5,7\}}$ equals $\sigma_2 \sigma_1 \sigma_7 \sigma_5 \sigma_4$. The transition $u \rightarrow \sigma_S u$ is given rate $\prod_{i \in S} x_{u_i}$.

Theorem 6.1. *Fix $k \in (0, n)$. The Markov chain Θ_k on words whose outgoing transitions from a word u are all $u \rightarrow \sigma_S u$, $S \subseteq [n]$, $|S| = k$, has the same stationary distribution for all k .*

In particular Θ_k has the same stationary distribution as Θ_1 for each k , where Θ_1 is the inhomogenous TASEP on a ring introduced in [7].

A weaker theorem has been proved by Martin and Schmidt [8], where the underlying graph is the infinite discrete line, all $x_i = 1$, and subsets S of size k are instead chosen at rate p^k , where $p \in (0, 1)$ is a parameter. Since our proof will be purely local, the theorem above implies their result.

The theorem can be reduced to proving that the transition matrices of all the Θ_k commute among themselves. To see this, note that the stationary distribution corresponds to the largest eigenvalue n and that the corresponding eigenspace is one-dimensional. In the next section we prove this commutation property.

Before that, let us note that in the homogenous case, $x_1 = \dots = 1$, the chain M_k in Theorem 6.1 has a purely geometric definition: randomly compose a subset S of k simple generators (together with the reflection in the highest root), ordering them according to some orientation of the affine Dynkin diagram (the graph with




FIGURE 4. A diagram D . The set $C(D)$ is given by all words $u_1 \dots u_6$ such that u_2 is larger than all other letters u_1, u_3, u_4, u_5, u_6 , and $u_4 < u_3 < u_5$. For example, the crossing down to the left implies that $u_3 > u_4$.

one node for each generator and one node for the highest root, where two nodes are joined by an edge if their order is > 2 (i.e. if they do not commute)). We have tried all possible orientations of the affine Dynkin graph (on 4 nodes) of the group B_3 , but none is consistent with a result analogous to Theorem 6.1.

Finally, we remark that the operators σ_S are closely related to the multi-line queues [4]. Indeed, a special case of the case $k = n - 1$ of the theorem follows directly from the theory of multi-line queues: for readers familiar with multi-line queues, this is given by comparing the $(n - 1)$ -TASEP with multiline queues which have an extra last row to which no new particles are added. Thus, Theorem 6.1 interpolates between the definition of the chain ($k = 1$) and its highly non-trivial description in terms of multi-line queues ($k = n - 1$).

6.1. Proof of Theorem 6.1. As noted earlier, it suffices to show that for any $k, l \in (0, n)$, the transition matrices A_k, A_l of M_k and M_l commute.

For words u, v , the (u, v) entry of $A_k A_l$ is a weighted count of pairs (S, T) such that $v = \sigma_T \sigma_S u$, $|S| = l$ and $|T| = k$ (weighted by the product of the rates of the two transitions $u \rightarrow \sigma_S u \rightarrow \sigma_T \sigma_S u$). Similarly, the entry $(A_l A_k)(u, v)$ counts pairs (S', T') such that $v = \sigma_{T'} \sigma_{S'} u$, $|S'| = k$, $|T'| = l$. The idea of the proof is to find a weight-preserving involution from pairs (S, T) to (S', T') . This turns out to be tricky. Without loss of generality, we will only consider words with distinct letters – the statement for general words follows by merging particle classes.

Given a triple (u, S, T) , define a $2 \times n$ array D (cyclic in the horizontal direction) as follows. If $i \in S$ and $u_i < u_{i-1}$, color the site $(1, i)$ in the array black (\bullet). Otherwise, if $u_i > u_{i-1}$, color it white (\circ). Similarly, if $w = \sigma_S u$ and $i \in T$, color site $(2, i)$ black if $w_i < w_{i-1}$ and white if $w_i > w_{i-1}$. We refer to a site which is not colored as *empty*. Any $2 \times n$ array which arises in this way will be called a *diagram*. Let $C(D)$ denote the set of words u which (together with some sets S, T) give D as above. An example of a diagram is given in Figure 4, where its associated *particle trajectories* are also given. Note that these lines may be added or omitted as we please – they are determined by the coloring of the $2 \times n$ array. Formally, a trajectory is given by the three positions p_1, p_2, p_3 of the particle in $u, \sigma_S u, \sigma_T \sigma_S u$ respectively. We then say that the particle has *visited* position p_1 in the upper row and position p_2 in the lower row (we will use no such notation for p_3). Instead of defining an involution on triples (u, S, T) , we will define an involution on diagrams. Though diagrams are considered cyclic, we will most often deal with segments of diagrams.

Thus we can think of a diagram as describing the trajectories of the particles u_1, \dots, u_n .




FIGURE 5. Another diagram.

Say that two diagrams D, D' of the same length n are *compatible* if the following conditions hold.


- (1) The order relations implied by D and D' are equivalent. That is, $C(D) = C(D')$.
- (2) The number of colored sites in the top row of D equals the number of colored sites in the bottom row of D' and conversely.
- (3) The number of black sites passed by particle i is the same in D as in D' , for each i , and the same holds for white sites.

As an example, the diagrams in Figures 4 and 5 are compatible. Not all assignments of $\{\text{white, black, empty}\}$ to the sites of a $2 \times n$ array gives a diagram - for example, the $2 \times n$ array with only two colored sites, one white atop a black one is not a diagram (if a particle failed to jump the first time and the particle in front of it did not move, it will fail again).

To prove the theorem, it suffices to construct an involution on diagrams pairing up compatible diagrams - condition (3) guarantees that two compatible diagrams have the same weight (for any input word $u \in S(D) = S(D')$).

We will construct the involution α by successively restricting the set of all diagrams to yet smaller classes of diagrams, and show how to 'lift' any definition of an involution on a smaller class to a bigger one. These restrictions all involve finding some subdiagram and replacing it by a smaller one. They are defined as follows.

- (1) For a diagram D with the following sub- 2×2 -diagram in column c and $c + 1$,




let $\varphi^{(c)}(D)$ denote the diagram where the sub- 2×2 -diagram is replaced by the following 2×1 -diagram (in a single new column replacing $c, c + 1$).



Here, A, B are placeholders for any of \bullet, \circ , or an empty space.

- (2) (a) If D contains the 2×2 subdiagram




in column c and $c + 1$, let $\varphi^{(c)}(D)$ denote the result of replacing the 2×2 diagram by the 2×1 diagram



A

(b) If D is a diagram containing the 2×2 subdiagram





in column c and $c + 1$, let $\varphi_{\circ}^{(c)}(D)$ denote the result of replacing the 2×2 diagram by the 2×1 diagram

A



(3) If D contains the sub- 2×2 -diagram



in column c and $c + 1$, let $\varphi_{\circ}^{(c)}(D)$ denote the diagram with the subdiagram replaced by the following.

A

B

To find $\alpha(D)$ for a diagram D , we first perform three kinds of reductions. After these reductions we make a small change to the diagram, and then invert the reductions. The whole process is illustrated in Figure 6. The proofs that the reductions 'work' are all similar – we do all the details for one case in Lemma 6.3 and leave the remainder to the reader.

Lemma 6.2. *Let D, D' be two diagrams and c a column for which $\varphi_{\circ}^{(c)}$ is applicable. If $\varphi_{\circ}^{(c)}(D)$ and $\varphi_{\circ}^{(c)}(D')$ are compatible, then D and D' are compatible.*

Call a diagram D for which there is no c such that $\varphi_{\circ}^{(c)}$ can be applied a *I-reduced diagram*. The lemma shows that to construct our required involution it suffices to construct it on the set of I-reduced diagrams. An example of our involution α to be defined on this restricted set is given by removing the forbidden subdiagram from Figures 4 and 5.

It is easy to see that for a I-reduced diagram, two sites in the same column cannot both be black, and if there is a column colored black in the lower row and white in the upper row, then the site in the lower row to the left of the column is black.

Lemma 6.3. *Suppose D and D' are I-reduced diagrams, c a column. Then*

- *D and D' are compatible whenever $\varphi_{\circ}^{(c)}(D)$ and $\varphi_{\circ}^{(c)}(D')$ are both defined and compatible.*




FIGURE 6. Example of computing $\alpha(D)$ for the diagram D in the top. The three diagrams after the first one correspond to I-, II- and III-reductions. The next step consists of computing $\alpha(D')$ for the III-reduced diagram D' . The remaining steps consists of inverting the III-, II- and I-reductions. So the last diagram is the image under the involution of the first. To reduce clutter we have drawn the trajectories only for the first diagram.

- D and D' are compatible whenever $\varphi_{\bullet\bullet}^{(c)}(D)$ and $\varphi_{\bullet\bullet}^{(c)}(D')$ are both defined and compatible.
- D and D' are compatible whenever $\varphi_{\bullet\bullet}^{(c)}(D)$ and $\varphi_{\bullet\bullet}^{(c)}(D')$ are both defined and compatible.

Proof. We consider the second case. Denote the placeholders in D and D' by A and A' . Note that since the diagrams are I-reduced, A (and A') is not colored black. First we need to show that the output of D and D' are the same (for any input word $u \in C(D) = C(D')$), that is, we need to show that each particle ends up in the same place in both D and D' . We know this is true for $\varphi_{\bullet\bullet}^{(c)}(D)$ and $\varphi_{\bullet\bullet}^{(c)}(D')$ so we only need to consider particles that pass through column c in D . This is easy to check – if the particle passes column c in $\varphi_{\bullet\bullet}^{(c)}(D)$ it must do so in the upper row, and then it must pass column c in the lower row in $\varphi_{\bullet\bullet}^{(c)}(D')$. The column added when constructing D respectively D' preserves this property. And if the particle does not pass column c in $\varphi_{\bullet\bullet}^{(c)}(D)$ and $\varphi_{\bullet\bullet}^{(c)}(D')$, the same holds in D and D' . Finally it's easy to check that the particle starting in the 'new' column in D and D' end up in the same place (indeed, it's the only place left). Next, we need to show that $C(D) = C(D')$. This amounts to checking what new relations on the input word u the colors of the new sites in D and D' give us. This can be checked directly and individually for the new white site in D and D' and the new black site in D and D' . \square

Call a I-reduced diagram for which no $\varphi_{\bullet\bullet}^{(c)}$ nor $\varphi_{\bullet\bullet}^{(c)}$ is defined a II-reduced diagram. We now make the final restriction.

Lemma 6.4. *Suppose D, D' are II-reduced diagrams, and c is a column for which $\varphi_{\bullet\bullet}^{(c)}(D)$ and $\varphi_{\bullet\bullet}^{(c)}(D')$ are both defined. Then, if the latter two are compatible, then so are D and D' .*

If D is a II-reduced diagram for which no $\varphi_{\bullet\bullet}^{(c)}$ is defined we call D a III-reduced diagram. The variety of III-reduced diagrams is sufficiently narrow to be analyzed directly. Note that in any III-reduced diagram, if a particle passes two colored sites, then those are necessarily white. Suppose D is a III-reduced diagram. Above the starting column of each particle (see Figure 6), write either U, L or T , depending on whether the particle passes (T)wo sites colored white or a unique colored site in the (U)pper row or a unique colored site in the (L)ower row. To define $\alpha(D)$, for each maximal word $U^r D^s$ strictly between two words of the type T or LU , change the behavior of the particles corresponding to $U^r L^s$ so that it will read $U^s L^r$ instead (see the example in Figure 6). We define the resulting diagram to be $\alpha(D)$.

It is readily checked that the map is an involution, and that it pairs up compatible diagrams. Thus this is true of the extended involution on all diagrams, too. This finishes the proof.

7. QUESTIONS, REMARKS

7.1. Is there an easily described conjugation matrix $U_{i,J}$ for general links (i, J) in type C (or B or D)?

As an example, in the notation of Section 2, consider the link $(i, J, J') = (2, \{3, 4\}, \{2, 3, 4\})$ when $n = 4$. Convert states in Θ_J (and $\Theta_{J'}$) to words $w_1w_2w_3w_4$. If the i th column is empty, let $w_1 = 3$ (2), otherwise let $w_i = +i$ if there is a particle of class i in the upper line, and $w_i = -i$ if there is a particle of class i in the lower line. We order the 48 states of Θ_J lexicographically: $\bar{2}133 < \bar{2}133 < \bar{2}3\bar{1}3 < \dots < 3321$. Similarly, the 8 states of $\Theta_{J'}$ are ordered (from left to right) as $\bar{1}222 < 1222 < 2\bar{1}22 < 2122 < 22\bar{1}2 < 2212 < 222\bar{1} < 2221$. Indexing rows and columns this way, the transpose U of the following matrix satisfies $M_JU = UM_{J'}$.

$$\begin{pmatrix} 101021030100040200202042102102000100102142010021 \\ 101010030111040222202020101002110111101031121110 \\ 101121030000040000202242112101000000112143002100 \\ 101000030121040242202000100002210121100022212100 \\ 111121020000020000222242122100001121000043002100 \\ 100000031121042242200000010011211121000043002100 \\ 121121010000131121111121122100001121000043002100 \\ 010000121121242242000000122100001121000043002100 \end{pmatrix}$$

(This is a 8×48 matrix all of whose entries are in $\{0, 1, 2, 3, 4\}$.) Is there a combinatorial rule (along the lines of the queueing process of Section 3) which produces the column corresponding to any given state in $\Theta_{J'}$?

7.2. Is it possible to carry the explicit description of the stationary distribution Θ_J further than is done in Section 4, say for $|J| = n - 2$?

7.3. Can the matrix U in Section 3 be defined without reference to the permutation representation of the group, ie. using only the realization of the group as a reflection group?

7.4. Small examples indicate that a similar queueing process exists for groups of type B and D. Is it easier to extend the analysis in Sections 3 and 4 for these groups?

7.5. Can the k -TASEP be extended to general Weyl groups?

REFERENCES

- [1] Arvind Ayyer and Svante Linusson, An inhomogenous multispecies TASEP on a ring (2012). [arXiv:1206.0316](#)
- [2] Chikashi Arita, Arvind Ayyer, Kirone Mallick and Sylvain Prolhac, Recursive structures in the multispecies TASEP, *J. Phys. A* 44, 335004 (2011).
- [3] Anders Björner and Francesco Brenti, *Combinatorics of Coxeter Groups*, Springer-Verlag, New York, 2005.
- [4] Pablo Ferrari and James Martin, Stationary distributions of multi-type totally asymmetric exclusion processes, *Ann. Prob.* 35 (2007) 807–832.
- [5] B. Derrida, M. R. Evans, V. Hakim, and V. Pasquier, An exact solution of a 1D asymmetric exclusion model using a matrix formulation, *J. Phys. A* 26:1493–1517 (1993).
- [6] Thomas Lam, The shape of a random affine Weyl group element, and random core partitions (2011). [arXiv:1102.4405](#)
- [7] Thomas Lam and Lauren Williams, A Markov chain on the symmetric group which is Schubert positive?, *Experimental Mathematics* 21 (2012), 189–192.
- [8] James Martin and Philipp Schmidt, Multi-type TASEP in discrete time (2010). [arXiv:1002.3539](#)

DEPARTMENT OF MATHEMATICS, ROYAL INSTITUTE OF TECHNOLOGY, SE-100 44 STOCKHOLM, SWEDEN

E-mail address: eaas@kth.se